DETAILED ACTION

1. Claims 1-8, 10-36, 39-50, 52, 53, 57-63, 66, and 67 are pending.

Information Disclosure Statement

2. The IDS's dated May 15, 2008 and July 16, 2008 have been considered. See enclosed PTO-1449.

Allowable Subject Matter

3. Claims 35 and 36 are allowable over the prior art of record, however a Double Patenting Rejection is presented below.

4. The Examiner has indicated allowable subject matter which, if all the following limitations are incorporated into the other independent claims, would put the claims in condition for allowance (subject to the Double Patenting rejection presented below). The limitations are:

- Creating an unattached connection endpoint containing information associated with the client and the first intelligent network device; and
- Deleting the unattached connection endpoint in response to receiving an acknowledgement indicating that the second intelligent network device received the packet containing the unattached endpoint information.

Double Patenting

5. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory obviousness-type double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the conflicting application or patent either is shown to be commonly owned with this application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement.

Effective January 1, 1994, a registered attorney or agent of record may sign a terminal disclaimer. A terminal disclaimer signed by the assignee must fully comply with 37 CFR 3.73(b).

Claims 6-10, 16-26, 29-67 of Application no. 09/872,332 contains every element of claims 1-8, 10-36, 39-50, 52, 53, 57-63, 66, and 67 of the instant application and as such anticipates claims 1-8, 10-36, 39-50, 52, 53, 57-63, 66, and 67 of the instant application.

6. The Examiner has indicated allowable subject matter as presented above. Accordingly, the Examiner requests a Terminal Disclaimer in order to overcome the Double Patenting rejection above.

Claim Rejections - 35 USC § 112

7. The Office has considered the amendments to the claims. The rejection under this heading is withdrawn.

Claim Rejections - 35 USC § 103

8. The text of those sections of Title 35, U.S. Code not included in this action can be found in a prior Office action.

Claims 1-8, 10-34, 42-48, 53, 57-63, 66 and 67 are rejected under 35 U.S.C. 103(a) as being unpatentable over Aversa et al. (Load Balancing a Cluster of Web Servers Technical Report BUCS-TR-1999-01, Boston University, Computer

Science Department, January, 1999) (cited by Applicant in IDS) (hereinafter Aversa) in view of Snoeren et al. (TCP Connection Migration, Internet Draft, November 2000) (cited by Applicant in IDS) (hereinafter Snoeren) in view of Giles et al. (USPN 7,165,120) (hereinafter Giles).

9. Referring to claim 1, Aversa discloses an information processing system, comprising:

a first computing device (in the example given in the reference it is referred to as "Server 4") (Figure 2) for:

receiving an initialization packet (i.e. SYN) originating from a client (p. 5, ¶ 2-3);

storing information representing a connection with the client (i.e. inherent to TCP/IP that a TCB is created and maintained for the connection to track information regarding the session)

in response to at least the initialization packet, outputting a response packet (i.e. ACK) to the client (it is an inherent feature of the system that an ACK is sent to the client in order to tell the client that the connection is opened, this is one of the building blocks of the HTTP protocol);

receiving a request packet originating from the client (p. 4, ¶ 2); and

in response to at least the request packet and a state of at least one of the first computing device and a second computing device, selectively outputting a migration packet representing the connection with the client to the second computing

device and stores information representing the connection in a forward table (i.e. routing table)(p. 4, ¶ 2).

Aversa does not explicitly disclose removing the information representing the connection with the client from the temporary table in response to an acknowledgement indicating that the second computing device received the migration packet. Snoeren discloses a method for migrating TCP sessions which discloses sending a migration packet to another device, and upon acknowledgement, the connection is closed with the first device, essentially deleting the TCB from the first device (i.e. transfers the original TCB to the second connection) (p. 5, ¶ 2). It would have been obvious to combine the teaching of Snoeren with Aversa in order to provide an efficient method to transfer TCP connections to another device, resulting in increased efficiency and less wasted bandwidth.

Aversa-Bruck does not explicitly disclose that the first computing device conducts load balancing in a network interface device. In analogous art, Giles discloses another server network system which discloses a NIC processor 825 which performs load balancing/routing/switching, which performs these features without invoking the main computer (Figure 8; col. 5, line 50 to col. 6, line 2). It would have been obvious to one of ordinary skill in the art to combine the teaching of Giles with Aversa-Snoeren by replacing the software load balancing system described in Aversa with the NIC processor described in Giles in order to realize the benefits described in Giles, namely the inherent speed increases of hardware versus software implementation as well as providing a scalable server system to route requests between various server nodes.

Aversa-Snoeren-Giles

10. Referring to claim 2, Aversa-Snoeren-Giles discloses the first computing device is a NIC (this is an inherent feature of Aversa, since the Server would be unable to communicate with the network if a NIC was not installed on the Server) (Giles: Figure 8, ref 825).

11. Referring to claim 3, Aversa discloses the second computing device is configured to perform an operation of a software application (i.e. a web server) (p. 5, ¶ 3).

12. Referring to claim 4, Aversa discloses the software application is a socket based application (the Office takes the term "socket based application" as any application which either uses, controls, or interfaces with a socket, such as a web server) (p. 5, ¶ 3).

13. Referring to claim 5, Aversa discloses the initialization packet (i.e. SYN packet) is addressed by the client to the first computing device (i.e. Server 4), and wherein the first computing device is for receiving the initialization packet in response to the addressing (an inherent feature of HTTP is that a server always receives a packet to which it is addressed, barring any connection disruptions) (p. 4, ¶ 2).

14. Referring to claim 6, Aversa discloses outputting a response packet to the client and wherein the first computing device is for:

in response to at least the request packet and the state (i.e. server load), selectively outputting the request packet to the second computing device for outputting the response packet to the client, such that the output response packet bypasses the first computing device (i.e. Server 2, responds by sending the requested resource to the client using Server 4 as the source address) (p. 4, ¶ 2; p. 5, ¶ 3).

15. Referring to claim 7, Aversa discloses the first computing device is for receiving the initialization packet through a global computer network (i.e. Internet) (e.g. abstract).

16. Referring to claim 8, Aversa discloses the first computing device is for selectively outputting the migration packet to the second computing device through a LAN (Figure 2).

17. Referring to claim 10, Aversa discloses the information representing the connection with the client includes a group of sequence numbers associated with the connection (it is well known that TCB's include various sequential numbers to keep track of the connection such as window size, start sequence number, etc.).

18. Claims 11 and 13 are rejected for similar reasons as stated above. Furthermore Aversa discloses bypassing the first computing device and appearing to the client as

received from the first commutating device (i.e. Server 2 serves the request to client B using the IP address of Server 4 as the source address) (p. 3, ¶ 3).

19. Referring to claims 12 and 29, Aversa in view of Bruck discloses the invention substantively as described in claim 10. Aversa in view of Bruck does not specifically state the group of sequence numbers includes at least one start sequence number, at least one current sequence number, and at least one ACK number. It is well known that HTTP provides for storing current sequence numbers and ACK numbers in order to determine which packets have been received and acknowledged to determine if the requested resource has been fully received. In regards to the start sequence number, it would make sense to track that number for encryption purposes (some encryption algorithms start their sequence numbers at random numbers to confuse any snooping algorithms which think they have missed the beginning part of the sequence). By this rationale it would have been obvious to one of ordinary skill in the art to track the start sequence number, the current sequence number, and the ACK number to keep track of all the packets incoming and outgoing from the computing device, thereby providing for a reliable transport protocol for web documents.

20. Referring to claims 14-16, Aversa discloses the address includes an IP address and a TCP port, (i.e. an IP address of the first computing device and port 80, which is the port of the TCP/IP stack) (p. 5, ¶ 7).

21. Referring to claims 17 and 34, Aversa discloses the invention substantively as described in claim 15. Aversa further discloses the port is a TCP port as seen in claim 16, however remains silent over if the port can be a UDP port. It is well known that web servers can exist using UDP ports, and some applications rely upon these ports (web conferencing, gameplay, blizzard.net, etc) in order to provide an almost real-time feel to the content. It would have been obvious to one of ordinary skill in the art to provide for a UDP port to communicate the first computing device to a second computing device in order to rely upon the UDP protocol for a server, thereby increasing the abilities of the web servers and allowing more users access to the network.

22. Claims 18-28, 30-33, 42 and 43 are rejected for similar reasons as stated above.

23. Referring to claim 44, Aversa discloses the invention substantively as described in claim 42. Aversa does not specifically state the third packet is transmitted when network packet of the second type is received (Aversa discloses the synchronization information is transmitted “intermittently”). However one of ordinary skill would realize the benefits of transmitting the synchronization information when a packet has been received to ensure the updating of the load information for that particular server occurs when it is needed (i.e. when a new request arrives) and is reduced for when it is not needed (i.e. when a request does not come in for a period of time), thereby reducing overall traffic load on the network as well as reducing overall overhead processing of packets in the network.

24. Claims 45-48, 53, 57-63, 66 and 67 are rejected for similar reasons as stated above. Furthermore Bruck discloses a third network which transmits synchronization information (i.e. note lateral lines between servers 230 indicate state sharing and unique protocol words are communicated between the machines) (col. 6, line 45 to col. 7, line 2). Bruck further discloses the ports are physically separate from one another (i.e. a NIC interface for each network internal and external to which the machine is connected) (col. 7, lines 65-67).

Claims 49, 50 and 52 is rejected under 35 U.S.C. 103(a) as being unpatentable over Aversa-Snoeren-Giles as applied above in view of Joffe et al. (USPN 6,185,619) (hereinafter Joffe).

25. Referring to claim 49 and 52, Aversa discloses the invention substantively as described in claim 49. Aversa does not specifically disclose not using IP-IP encapsulation and without using TCP splicing techniques in order to output the packet to the first computer system. IN analogous art, Joffe discloses another information processing system which discloses not using IP-IP encapsulation and without using TCP splicing techniques in order to output the packet to the first computer system (col. 12, lines 50-55). It would have been obvious to one of ordinary skill in the art to combine the teaching of Joffe with Aversa in order to select an appropriate server from which to retrieve a data object for a user based upon the user's request and the

capabilities and topology of the underlying network as supported by Joffe (col. 3, lines 35-40).

26. Referring to claim 50, Aversa discloses the invention substantively as described in claim 49. Aversa does not specifically state sending a response packet by the first computer system such that the client thinks it came from the second computer system, however Aversa also states that the IP source addresses can be modified to include other servers as sending the response (see rejections above), and it is inherent that TCP packets receive ACK packets for sent packets, so therefore one of ordinary skill in the art would realize the benefits of having the client perceiving the response packet is received from the second computer system in order to remove the first computer system from repeatedly having to forward requests to the second computer system, thereby reducing the load on the first computer system (i.e. Server 4 still receives requests for Client B even though the request is being served by Server 2, however Server 4 still has to process the packet, determine which server is handling this request, and then forward this packet, costing valuable processing time) (p. 3, ¶ 2).

Claims 39-41 are rejected under 35 U.S.C. 103(a) as being unpatentable over Aversa-Snoeren-Giles as applied above in view of Brendel et al. (USPN 5,774,660).

27. Referring to claim 39, Aversa-Snoeren-Giles disclose the invention as described in the claims above. Aversa-Snoeren-Giles do not expressly disclose the receiving a

PUSH packet originating from the client, and then based on information transmitted between the servers outputting a migration packet, and encapsulating the PUSH packet to the first computing device of the second server. Brendel discloses a method of migrating TCP connections which shows the client establishing a session with the load balancer, receiving a PUSH packet, selecting a server, and sending the PUSH packet to the second server (Figure 11A). It would have been obvious to combine Brendel with Aversa-Snoeren-Giles in order to provide an efficient method to transfer connections based on what file is being requested.

28. Claims 40 and 41 are rejected for similar reasons as stated above.

Response to Arguments

29. Applicant's arguments dated May 15, 2008 have been fully considered but are moot in view of the new grounds of rejection presented above.

Conclusion

30. Applicant has failed to seasonably challenge the Examiner's assertions of well known subject matter in the previous Office action(s) pursuant to the requirements set forth under MPEP §2144.03. A "seasonable challenge" is an explicit demand for evidence set forth by Applicant in the next response. Accordingly, the claim limitations the Examiner considered as "well known" in the first Office action, are now established

as admitted prior art of record for the course of the prosecution. See *In re Chevenard*, 139 F.2d 71, 60 USPQ 239 (CCPA 1943).

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Joseph E. Avellino whose telephone number is (571) 272-3905. The examiner can normally be reached on Monday-Friday 7:00-4:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Nathan J. Flynn can be reached on (571) 272-1915. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Joseph E. Avellino/

Primary Examiner, Art Unit 2146